Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1.-19. (canceled)

20. (previously presented) A method for allocating bandwidth on a communication link in an operating network, comprising:

receiving a flow of data packets at a receiving system from a sending system on the network via the link;

determining at the receiving system a target bandwidth for the flow on the link;

transmitting data from the receiving system to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that when the sending system transmits subsequent data packets to the receiving system, such subsequent data packets are transmitted at a rate approximating the target bandwidth determined by the receiving system; and

calculating a period of time for which the receiving system does not transmit data to the sending system.

21. (previously presented) The method of claim 20 wherein a rate at which the transmitted data from the receiving system arrive at the sending system determines the rate at which the sending system transmits the subsequent data packets.

22. (previously presented) The method of claim 20, further comprising:

applying the calculated period of time for which the receiving system does not transmit data to a plurality of flows of a class.

23. (previously presented) The method of claim 20, wherein the data transmitted from the receiving system to the sending system indicates a maximum amount of data that the receiving

system will accept from the sending system in a subsequent data transmission.

24. (previously presented) The method of claim 20, wherein the step of calculating the period of time includes using a model of the network.

25. (previously presented) The method of claim 20, further comprising the steps of measuring a bandwidth of the link being used by a class of flows, and using the measured bandwidth as feedback for calculating the period of time.

26. (previously presented) The method of claim 20, wherein the data transmitted by the receiving system to the sending system includes acknowledgment of receipt of a particular data packet in the flow of data packets.

27. (previously presented) The method of claim 20, wherein the target bandwidth is determined by a network protocol process.

28. (previously presented) The method of claim 20, wherein when the link is idle, the target bandwidth equals a full bandwidth of the link.

29. (previously presented) A method for allocating bandwidth on a communication link in an operating network, comprising:

receiving flows of data packets at a receiving system from sending systems on the network via the link;

determining at the receiving system a target bandwidth for the flows on the link;

transmitting data to the sending systems that will cause the sending systems to transmit subsequent data packets to the receiving system at a rate approximating the target bandwidth for the flows of data packets from the sending systems; and

for each flow of at least a subset of the flows, calculating a period of time for which the receiving system does not transmit data to the sending system.

30. (previously presented) A computer system for allocating bandwidth on a communication link in an operating network, comprising:

a network interface for receiving a flow of data packets from a sending system on the network via the link; and

a processor coupled to the network interface for determining a target bandwidth for the flow on the link, the network interface transmitting data to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that the sending system transmits subsequent data packets at a rate approximating the target bandwidth for the flow when responding to the transmitted data;

wherein the processor is configured to calculate a period of time for which the receiving system does not transmit data to the sending system.

31. (previously presented) The computer system of claim 30, further comprising:

memory, coupled to the processor, configured to store a model of a behavior of the network that is used when calculating a period of time for which the receiving system does not transmit data to the sending system.

32. (previously presented) The computer system of claim 30, wherein the processor is configured to assign the flow to a class of flows; and wherein a calculated period of time for which the receiving system does not transmit data to the sending system applies to the flows in said class.

33. (previously presented) The computer system of claim 30, wherein the system is configured to measure a bandwidth of the link being used by a class of flows and to calculate, based on the measured bandwidth, a period of time for which the receiving system does not transmit data to the sending system.

34. (previously presented) The computer system of claim 30, wherein the data transmitted from the receiving system to the sending system include a maximum amount of data that the receiving system will accept from the sending system in a subsequent data transmission.

35. (previously presented) The computer system of claim 30, wherein the target bandwidth is determined by an application program that receives the flow of data packets.

36. (previously presented) A method for allocating bandwidth on a communication link in an operating network, comprising:

receiving a flow of data packets at a receiving system from a sending system on the network via the link;

determining at the receiving system a target bandwidth for the flow on the link; and

transmitting data from the receiving system to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that when the sending system transmits subsequent data packets to the receiving system, such subsequent data packets are transmitted at a rate approximating the target bandwidth determined by the receiving system; and

calculating and applying a period of time for which the receiving system does not transmit data to a plurality of flows of a class;

wherein the target bandwidth is determined by an application program that receives the flow of data packets.

37. (previously presented) The method of claim 36, wherein the step of calculating the period of time includes using a model of the network.

38. (previously presented) The method of claim 36, further comprising the steps of measuring a bandwidth of the link being used by a class of flows, and using the measured bandwidth as feedback for calculating the period of time.

39. (previously presented) A method for allocating bandwidth on a communication link in an operating network, comprising:

receiving a flow of data packets at a receiving system from a sending system on the network via the link;

determining at the receiving system a target bandwidth for the flow on the link; and

transmitting data from the receiving system to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that when the sending system transmits subsequent data packets to the receiving system, such subsequent data packets are transmitted at a rate approximating the target bandwidth determined by the receiving system;

wherein the target bandwidth is determined by an application program that receives the flow of data packets; and the data transmitted by the receiving system to the sending system includes acknowledgment of receipt of a particular data packet in the flow of data packets.

40. (previously presented) In a data flow control device operative to control the rate of data packets transmitted between first and second transmission stations in a packet communications environment, wherein the first transmission station is operative to transmit at least one packet associated with a data flow to the second transmission station, wait for acknowledgment of at least one transmitted packet before transmitting subsequent packets associated with the flow, and retransmit the at least one packet if an acknowledgment is not received with a period of time, a method comprising

delaying transmission acknowledgment packets, corresponding to a first data flow transmitted from the first transmission station to the second transmission station, by a computed delay, wherein the computed delay associated with the delayed acknowledgment packet provides feedback to the first transmission station such that when the first transmission station transmits subsequent packets corresponding to the flow, such subsequent packets are transmitted at a rate approximating a target rate;

receiving a second data flow from the first transmission station, wherein the data flow comprises at least one packet;

storing the data flow in a memory;

if the second data flow is a retransmission of the first data flow, determining whether an acknowledgement packet corresponding to the first data flow is in the memory; and

if so, deleting the second data flow from the memory; otherwise, forwarding the second data flow to the second transmission station.

41. (previously presented) The method of claim 40 further comprising classifying the first data flow into one of a plurality of classes to determine the target rate for the first data flow.

42. (new) A method for controlling the rate of data packets transmitted between first and second nodes in a packet communication environment, wherein the first node is operative to transmit at least one packet associated with a flow to the second node, and wait for acknowledgment of at least one transmitted packet before transmitting subsequent packets associated with the flow, said method comprising:

forwarding at least one packet corresponding to a flow from a first node to a second node;

receiving an acknowledgment packet from the second node to the first node, the acknowledgment packet acknowledging at least one packet in the flow transmitted from the first node, wherein the acknowledgment packet received from the second node includes a window size indicator that specifies an allowable range of transmission of data beyond a range of data acknowledged as a window size to be advertised from said second node to said first node;

selecting a substitute window size indicator for said window size indicator to modify the rate of transmission of packets from the first node;

inserting said substitute window size indicator into said acknowledgment packet; and

forwarding the acknowledgment packet to the first node.

43. (new) The method of claim 42 wherein the substitute window size indicator is computed based on a target rate for the flow.

44. (new) The method of claim 42 wherein the first and second nodes implement the TCP protocol and wherein the flow is a TCP flow.

45. (new) The method of claim 42 further comprising

computing a delay for transmission of the acknowledgment packet to control the rate of transmission of packets from the first node; and

delaying the forwarding step by the computed delay.

46. (new) The method of claim 42 wherein the substitute window size indicator is computed based on a one metric characterizing an attribute of the communications path between the first and second node.

47. (new) The method of claim 46 wherein the metric is the measured data rate between the first node and the second node.

48. (new) The method of claim 46 wherein the metric is the round trip time between the first node and the second node.

49. (new) A method for controlling the rate of data packets transmitted between first and second nodes in a packet communication environment, wherein the first node is operative to transmit at least one packet associated with a flow to the second node, and wait for acknowledgment of at least one transmitted packet before transmitting subsequent packets associated with the flow, said method comprising:

forwarding packets corresponding to a flow from a first node to a second node;

receiving acknowledgment packets from the second node to the first node, the acknowledgment packets acknowledging at least one packet in the flow transmitted from the first node, wherein the acknowledgment packets received from the second node includes a window size indicator that specifies an allowable range of transmission of data beyond a range of data acknowledged as a window size to be advertised from said second node to said first node;

computing, for at least one acknowledgement packet, a substitute window size indicator for said window size indicator to control the rate of transmission of packets from the first node;

inserting said substitute window size indicator into said acknowledgment packet; and

forwarding the acknowledgment packet, modified in the inserting step, to the first node.